Therapeutic Psychedelic LSD Analogues Without Hallucinations: JRT, Non-Hallucinogenic Neuroplastogens, and the Medicinal Chemistry of Safer Psychoplastogens

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

January 11, 2026

Neuropsychiatry is confronting a strange design challenge: many of the most potent “reset” effects on mood and addiction have been linked to compounds that also produce intense perceptual alterations. A new wave of medicinal chemistry tries to separate these two properties by redesigning psychedelic scaffolds into analogues that preserve neuroplasticity and therapeutic-like behavioral outcomes while reducing hallucinogenic potential. This paper uses the LSD analogue JRT as a focal case study because it captures the core idea in a single, memorable move: small structural edits can cause large changes in signaling and in animal proxy measures tied to psychedelic effects. From there, we map the broader landscape of “neuroplastogens” and “psychoplastogens” engineered from classic psychedelic families (ergolines, tryptamines, and iboga alkaloids), including candidates often described as non-hallucinogenic in preclinical assays. We also address the central interpretive hazard: “non-hallucinogenic” is frequently a statement about animal proxies and molecular signatures, not human subjective experience. The goal is a practical framework for readers who want to understand what is being optimized (receptor engagement, biased signaling, plasticity, safety margins), what counts as evidence at each stage, and how to talk about these compounds rigorously without hype.

psychedelic analogues, LSD analogue, JRT, non-hallucinogenic, reduced hallucinogenic potential, neuroplastogens, psychoplastogens, therapeutic psychedelics, 5-HT2A, biased agonism, neuroplasticity, antidepressant-like effects, addiction pharmacology, medicinal chemistry, structure-activity relationship, head-twitch response, translational biomarkers, safety pharmacology, ergoline, ibogaine analogue, tabernanthalog, 2-bromo-LSD, lisuride, zalsupindole, Delix Therapeutics.

A collaboration with GPT-5.2-Thinking. 50 pages. CC4.0.

Outline

1. Introduction: Why “Therapeutic Without Trip” Became a Design Goal
1.1 The clinical motivation: depression, addiction, and rapid-acting interventions
1.2 The scientific motivation: separating plasticity from perception
1.3 The communication problem: what “non-hallucinogenic” can and cannot mean

2. Terms and Operating Definitions
2.1 Psychedelics, analogues, and scaffold families
2.2 Neuroplastogens and psychoplastogens as functional categories
2.3 Hallucination proxies in animals vs human subjective experience
2.4 What counts as “promise”: preclinical signal, mechanistic clarity, translational plausibility

3. Case Study: JRT as an LSD Analogue Designed for Reduced Hallucinogenic Potential
3.1 The design premise: minimal structural change, maximal functional shift
3.2 Synthesis overview at the concept level: why de novo matters
3.3 Cellular signaling profile: what was optimized and what changed
3.4 Mouse behavioral proxies: interpreting “reduced hallucinogenic potential”
3.5 Antidepressant-like and cognitive-relevant assays: what they suggest, what they do not
3.6 The core lesson from JRT: small edits can re-route receptor-state landscapes

4. The Mechanistic Hypothesis Space
4.1 5-HT2A engagement as a lever, not a guarantee
4.2 Biased agonism and downstream pathway selection
4.3 Plasticity as a multi-node phenomenon: spines, synapses, circuits
4.4 Psychosis-relevant signatures vs therapeutic-relevant signatures
4.5 Dose, exposure, and kinetics: why timing can mimic “mechanism”

5. The Assay Stack: How the Field Decides a Compound Is “Less Psychedelic”
5.1 Behavioral proxies (for example, head-twitch response): strengths and limits
5.2 Gene-expression and circuit readouts: what they add beyond behavior
5.3 Receptor profiling and functional selectivity: common pitfalls
5.4 Safety pharmacology: cardiotoxicity, vasoconstriction, seizure risk, abuse liability
5.5 A proposed evidence ladder: from receptor to rodent to human

6. Other Analogues Showing Promise
6.1 Tabernanthalog (TBG): an ibogaine-inspired neuroplastogen
6.2 AAZ-A-154 and related “psychLight”-screened candidates
6.3 2-bromo-LSD: an ergoline line with recurring clinical interest
6.4 Lisuride as a comparator and the controversy over “non-psychedelic” labels

6.5 Zalsupindole and modern industrial pipelines for non-hallucinogenic neuroplastogens
6.6 Why these candidates matter as a set: convergences and divergences

7. Design Principles Emerging Across Scaffolds
7.1 Minimize perceptual liability while preserving plasticity: what “separation” looks like
7.2 Structural strategies: atom swaps, ring edits, substituent tuning, stereochemistry
7.3 Signaling strategies: partial agonism, functional selectivity, pathway bias
7.4 Systems strategies: kinetics, brain penetration, metabolite control
7.5 A practical medicinal chemistry checklist for “safer psychoplastogens”

8. Translational Strategy: From Animal Proxies to Human Claims
8.1 Biomarkers that might bridge species (EEG, imaging, peripheral readouts)
8.2 Trial design issues: expectation effects, blinding, and functional outcomes
8.3 Indication matching: depression vs addiction vs cognition
8.4 The regulatory reality: defining benefit without relying on subjective “trip” endpoints

9. Risks, Ethics, and Public Language
9.1 Hype control: avoiding the false binary of “hallucinogenic vs therapeutic”
9.2 Misuse risk and access: what changes when compounds feel “ordinary”
9.3 Equity and commercialization: patents, startups, and the ownership of scaffolds
9.4 A disciplined vocabulary for journalists, investors, and clinicians

10. Conclusion: The Safer-Psychoplastogen Thesis
10.1 What JRT contributes as a design exemplar
10.2 What the analogue landscape implies about mechanism
10.3 The near-term research agenda: decisive experiments and failure modes
10.4 A constructive closing: rigor as the path to trust and clinical legitimacy

11. References
11.1 Primary papers on JRT and LSD-analogue design
11.2 Key neuroplastogen and psychoplastogen discovery papers
11.3 Receptor signaling, biased agonism, and 5-HT2A mechanistic literature
11.4 Preclinical proxy assay validation and limitations
11.5 Clinical and regulatory perspectives on psychedelic-inspired therapeutics

Art

1. Introduction: Why “Therapeutic Without Trip” Became a Design Goal

The modern psychedelic revival has produced a paradox that is easy to state and hard to resolve: some of the strongest signals for rapid relief in depression, trauma-related suffering, and addiction have emerged around compounds that also produce intense alterations in perception, self-modeling, and meaning. For some patients, those experiences are central to healing. For others, they are frightening, destabilizing, medically risky, socially impractical, or simply incompatible with daily life and employment. As the field has matured, it has become clear that “psychedelic medicine” is not a single product category but a design space. One corner of that space is experience-centered therapy (the “trip” is part of the intervention). Another corner is experience-minimized pharmacology (the “trip” is treated as a side effect to be reduced). The phrase “therapeutic without trip” names the second corner: a push toward compounds that retain durable clinical benefit while reducing hallucination-like effects, dissociation-like effects, or other forms of profound altered consciousness.

This design goal is not merely a concession to stigma or comfort. It arises from practical constraints in health systems: scalable treatment models, cost and staffing, risk management, and equitable access. If an intervention requires many hours of supervised administration and specialized psychological support, it may remain boutique even if it is highly effective. Conversely, if medicinal chemistry can produce agents that deliver some of the same downstream therapeutic improvements with minimal disruption, the treatment could become far more scalable, and potentially reach patients who would never choose or tolerate a classic psychedelic session. The rise of engineered analogues is, in that sense, an attempt to translate a compelling but high-friction phenomenon into a lower-friction medicine.

1.1 The clinical motivation: depression, addiction, and rapid-acting interventions

The clinical motivation begins with unmet need. Major depression remains common, recurrent, and often resistant to first-line pharmacotherapy. Many people cycle through multiple antidepressants with partial relief, unacceptable side effects, or delayed onset that feels intolerable during crisis. Traditional medications often require daily adherence and weeks to demonstrate benefit, while symptom burden can be immediate and severe. The appeal of psychedelic-related therapies has been their repeated association with rapid shifts: changes in mood, rumination, and behavioral flexibility that occur on shorter timescales than conventional pharmacology. Even when the evidence base is still evolving, the clinical imagination has been captured by the possibility of a “rapid-acting” intervention that is not simply sedation, stimulation, or short-lived euphoria.

Addiction raises the stakes further. Substance use disorders commonly involve entrenched habits, strong cue reactivity, and relapse loops that resist purely cognitive approaches. Many

patients describe addiction as a narrowing of available action, a contraction of future imagination. Psychedelic-assisted paradigms have been framed as opening a window in which rigid narratives loosen and motivational structure can be reset. Yet addiction medicine also faces real-world constraints: patients may have unstable housing, complex medical comorbidities, legal exposure, and limited tolerance for prolonged clinical sessions. A compound that could promote durable behavioral change without requiring a long supervised psychedelic experience would be attractive to clinicians and health systems alike.

There is also a third, quieter motivator: heterogeneity. Not every patient wants a mystical or destabilizing experience, and not every patient can safely have one. A history of psychosis, bipolar vulnerability, severe dissociation, or certain cardiovascular risks can shift the risk-benefit equation. Even when classic psychedelics are not contraindicated, they may be undesirable for patients whose lives do not allow an all-day treatment session with recovery time. The "therapeutic without trip" program is thus partly about expanding eligibility. It is an attempt to preserve an emerging therapeutic class while widening the doorway into it.

Finally, there is an operational motivation. If a therapy is effective but expensive and labor-intensive, it may remain inaccessible. If it can be delivered more like standard pharmacotherapy, it might scale. Some of the analogue work can be read as a bid to move from an artisanal model toward something closer to routine care, while still preserving the distinctive "reset" signature that has drawn attention to psychedelics in the first place.

1.2 The scientific motivation: separating plasticity from perception

The scientific motivation is more ambitious: it asserts that the therapeutic value of psychedelic-class compounds may be separable from their most conspicuous phenomenology. This is not a settled claim; it is a hypothesis that has become a guiding question. What, exactly, is the active ingredient? Is it the subjective experience itself (insight, meaning, fear resolution, self-transcendence)? Is it a biological cascade (synaptic remodeling, altered network dynamics, gene expression changes) that can occur with or without dramatic perceptual change? Or is it the interaction between the two: biology creates a window, experience and context steer what happens during that window?

The "separation" program tends to emphasize plasticity. In this framing, the brain is not merely "chemically corrected," but temporarily shifted into a state of elevated malleability. That malleability could allow patients to rewire stuck patterns, integrate learning, and rebuild adaptive behavior. If that is the core, then the obvious next question is whether the same plasticity can be triggered with less disruption to consciousness.

This is where receptor pharmacology and signaling nuance enter. Classic psychedelics strongly engage serotonin 5-HT2A signaling, but receptors are not binary switches. The same receptor

can couple into multiple downstream pathways depending on ligand structure, receptor conformation, cellular context, and kinetics. “Biased agonism” and functional selectivity have become central concepts because they offer a plausible mechanism for partial separation: two ligands can bind the same receptor yet emphasize different downstream cascades. An analogue might preserve signaling that promotes synaptic growth while reducing signaling patterns associated with hallucination-like effects.

Another scientific motivator is the recognition that “psychedelic effect” is not one thing. Perception, self-processing, emotional salience, and narrative meaning are mediated by different circuits and different timescales. Medicinal chemistry can potentially tune onset, duration, brain penetration, and metabolite profiles, all of which shape experience. In other words, “trip intensity” is not only about which receptor is hit; it is also about how quickly, how long, and in what brain compartments the system is perturbed. Analogue programs therefore mix structural design (what the ligand is) with kinetic design (how the ligand behaves in the body).

The JRT story is compelling because it dramatizes the idea in a minimalist way: a small structural edit can re-route functional outcomes. But the deeper scientific point is general. The field is testing whether therapeutic leverage can be extracted from the psychedelic family by moving from “phenomenology-first” to “mechanism-and-outcome-first,” without denying that experience can matter for some patients.

1.3 The communication problem: what “non-hallucinogenic” can and cannot mean

This entire program lives or dies on language. The public phrase “non-hallucinogenic psychedelic” is rhetorically powerful and scientifically treacherous. In most cases, it does not mean what an ordinary reader assumes it means. It usually means “non-hallucinogenic in preclinical proxy assays” or “reduced hallucinogenic potential according to animal behaviors and molecular signatures.” Those are not trivial claims, but they are not equivalent to “will not produce altered perception in humans.”

The communication risk has several layers. First, animal models cannot report subjective experience, so researchers use proxies. Some proxies correlate with certain receptor engagements and with the presence of classic psychedelics, but correlation is not identity. A compound can fail to trigger a proxy and still alter human consciousness, or trigger a proxy while producing an experience that is qualitatively different from what people imagine. Proxies are tools, not verdicts.

Second, the term “hallucination” itself is medically loaded. Clinical hallucinations in psychotic disorders are not the same phenomenon as the controlled perceptual alterations of psychedelic states. Collapsing these into one word can distort both psychiatry and psychedelic science.

When researchers say “reduced hallucinogenic potential,” they are often attempting to describe the reduction of psychedelic-like behavioral signatures, not claiming that they have eliminated psychosis risk or eliminated all altered states.

Third, overconfident language can provoke backlash. If early-stage compounds are marketed, even informally, as “no-trip psychedelics,” and later produce subjective effects in humans, trust erodes. Conversely, if the field insists that the “trip” is always essential, it may unnecessarily narrow therapeutic innovation and deny the real diversity of patient needs. The right stance is disciplined humility: describe what was measured, what was not measured, and what the measurements can support.

For this paper, “therapeutic without trip” is treated as an engineering objective, not a claim of final success. “Non-hallucinogenic” should be used, if at all, as shorthand for a specific kind of evidence: reduced psychedelic-like signals in established preclinical assays and reduced engagement of pathways thought to correlate with psychedelic phenomenology. Human experience remains an empirical question. The scientific task is to build a ladder from molecules to mechanisms to measurable outcomes, while the ethical task is to speak about each rung without pretending it is the top.

2. Terms and Operating Definitions

A field that moves fast tends to mint words faster than it earns stable meanings. Psychedelic-inspired drug discovery is exactly that kind of field: part medicinal chemistry, part clinical psychiatry, part cultural story. The result is a vocabulary that often slides between technical and popular uses in the same sentence. This section fixes operating definitions for the remainder of the paper. The goal is not to police language, but to prevent category errors—especially the common error of treating preclinical shorthand (“non-hallucinogenic”) as if it were a clinical conclusion. If the discipline is clear here, later sections can stay rigorous without becoming pedantic.

2.1 Psychedelics, analogues, and scaffold families

For purposes of this paper, a “psychedelic” is defined in the narrow pharmacological sense: a compound whose canonical effects in humans include marked alterations in perception, cognition, emotion, and self-processing, and for which the serotonin 5-HT2A receptor is a major mechanistic driver of those effects. This definition intentionally centers mechanism and phenotype rather than cultural status or legal category. It also excludes compounds sometimes grouped with psychedelics in popular discourse but driven primarily by other mechanisms (for example, dissociatives or deliriants), unless explicitly stated.

An “analogue” is defined structurally: a compound that is intentionally designed by modifying a known scaffold—often a classic psychedelic scaffold—to alter functional properties (efficacy, bias, kinetics, safety) while preserving enough structural continuity to learn from structure–activity relationships. The word “analogue” is not a promise of similarity in experience. It is a statement about design lineage and chemical resemblance. In other words, “LSD analogue” means “a molecule that lives near LSD in chemical space,” not “a molecule that feels like LSD.”

A “scaffold” is the core structural framework that chemists treat as the basis for a family of modifications. In psychedelic-inspired discovery, several scaffold families are frequently discussed:

- Ergoline or lysergamide family (LSD-like): polycyclic ergot-derived frameworks with amide substitutions that strongly influence potency, receptor residence time, and downstream effects.
- Tryptamine family (psilocin/psilocybin-like and DMT-like): indole-containing frameworks with substitution patterns that tune receptor profiles, half-life, and subjective intensity.
- Phenethylamine family (mescaline/2C-x/DOx-like): phenethylamine cores where ring substitutions and alpha substitutions shape duration, potency, and receptor selectivity.
- Iboga alkaloid family (ibogaine-like): complex polycyclic alkaloids with multi-receptor pharmacology, often pursued in addiction contexts; analogues here may pursue reduced toxicity and improved safety margins.

These families are useful because they give a map: when the field says “analogue,” the first question is “analogue of what scaffold family?” The second question is “what feature was preserved?” (binding, plasticity, duration) and “what feature was targeted for reduction?” (hallucination-like proxies, cardiotoxicity, vasoconstriction, abuse liability).

Finally, we distinguish “psychedelic-inspired therapeutics” from “psychedelics.” A psychedelic-inspired therapeutic may include classic psychedelics, but the phrase also covers engineered analogues and mechanistically adjacent compounds whose design motivation came from psychedelic biology, even if the final molecule is intended to be experience-minimized.

2.2 Neuroplastogens and psychoplastogens as functional categories

Two terms—neuroplastogen and psychoplastogen—appear frequently in this literature and are often used interchangeably. This paper treats them as related but not identical.

A “neuroplastogen” is defined operationally: any compound that robustly promotes measurable markers of neural plasticity, such as dendritic spine growth, synaptogenesis-related signaling, or circuit-level restoration, in ways that plausibly support therapeutic change. The key is function:

a neuroplastogen is whatever causes plasticity, regardless of how it feels, and regardless of whether it is culturally labeled a psychedelic. Under this definition, different mechanistic classes could qualify if they reliably promote plasticity.

A “psychoplastogen” is a narrower category: a compound that promotes neural plasticity in a way that is plausibly relevant to psychiatric symptom improvement, especially in mood and addiction domains, and that can produce enduring behavioral changes after transient exposure. The “psycho-” prefix highlights psychiatric relevance, not subjective psychedelic experience. In practice, many psychoplastogens have been serotonergic, but the definition does not require it.

Because these are functional categories, a key implication follows: calling something a psychoplastogen is not the same as calling it a psychedelic, and it is not the same as calling it non-hallucinogenic. It is a claim about a biological and behavioral package: plasticity plus durable functional outcomes in the kinds of assays that map (imperfectly) to clinical goals.

This distinction matters because it changes the reader’s default inference. If you read “psychedelic analogue,” you might expect “psychedelic-like experience.” If you read “psychoplastogen,” you should expect “plasticity-promoting therapeutic intent,” with experience as a separate question.

For this paper, “therapeutic without trip” candidates are described as psychedelic-inspired psychoplastogens that attempt to preserve plasticity and therapeutic-like behavioral outcomes while minimizing psychedelic-like signals in preclinical assays and, ideally, minimizing subjective psychedelic intensity in humans when that is later tested.

2.3 Hallucination proxies in animals vs human subjective experience

The most important operating definition in the paper is this: in preclinical work, “hallucinogenic” is almost never directly measured. It is inferred from proxies.

Because animals cannot report subjective hallucinations, researchers rely on behavioral and physiological signatures that correlate with known psychedelics. One of the most widely used is the head-twitch response in mice, a rapid stereotyped movement that correlates with 5-HT2A activation for many compounds. Other proxies include patterns of gene expression, immediate early gene activation, electrophysiological signatures, and certain behavioral paradigms that reflect altered sensory processing or cognitive flexibility.

These proxies are useful, but they can be misread. In this paper:

- A “hallucination proxy” is defined as an animal measure that correlates with classic psychedelics and is commonly used as an indicator of psychedelic-like activity.

- “Reduced hallucinogenic potential” means a compound produces a smaller signal on one or more of these proxies, especially relative to a benchmark psychedelic at behaviorally relevant doses.

What it does not mean is “no altered consciousness in humans.” It cannot mean that, because the mapping from proxy to subjective report is partial and context-dependent. The best way to state the relationship is: proxies are indicators of certain receptor-driven and circuit-driven effects that are often associated with psychedelic phenomenology in humans, but they are neither necessary nor sufficient conditions for the phenomenology.

Two additional problems complicate communication:

1. Psychedelic experience is multidimensional. A compound might reduce perceptual distortions while still altering affect, salience, time perception, or self-processing. If a social media post equates “no head twitch” with “no subjective alteration,” it is compressing a complex state into a single metric.
2. The word “hallucination” is clinically ambiguous. Psychedelic states are not identical to psychotic hallucinations, and many clinicians prefer different language (for example, “altered perception” or “psychedelic effects”) when describing classic psychedelics. Nevertheless, “hallucinogenic potential” persists because it is historically entrenched. When used here, it is treated as shorthand for “psychedelic-like proxy activity,” not as a claim about psychosis.

Therefore, the paper uses three distinct phrases that should not be collapsed:

- “Non-hallucinogenic in animal proxies” (a preclinical description)
- “Reduced psychedelic-like proxy activity” (a relative preclinical description)
- “Non-psychedelic or experience-minimized in humans” (a clinical description requiring human data)

If a compound is at the preclinical stage, only the first two are defensible, and even then only with explicit mention of what was measured.

2.4 What counts as “promise”: preclinical signal, mechanistic clarity, translational plausibility

The word “promise” is often used as a rhetorical placeholder for enthusiasm. Here we operationalize it as a three-part test. A candidate analogue “shows promise” only insofar as it clears three hurdles, each with its own evidence type.

First, preclinical signal. This means the compound produces robust, replicable effects in at least one well-justified preclinical assay relevant to the intended therapeutic domain. For mood, this

might include antidepressant-like behavioral assays plus evidence of plasticity; for addiction, it might include reduced drug seeking or relapse-related behaviors. “Robust” means dose–response relationships, appropriate controls, and ideally some replication across models or labs. A single striking finding can be interesting, but it is not strong “promise” unless it holds under variation.

Second, mechanistic clarity. This does not require full mechanistic understanding, but it does require more than hand-waving. At minimum, there should be credible linkage between molecular target engagement and downstream outcomes: receptor binding and functional assays that plausibly explain why the compound differs from its parent scaffold. Mechanistic clarity matters because it allows learning. Without it, each analogue becomes an isolated anecdote. With it, structure–activity relationships become a compass.

Third, translational plausibility. This is the hardest and most easily inflated. Translational plausibility is the argument that the preclinical package could become a workable human therapy. It includes pharmacokinetics (brain exposure, half-life, oral bioavailability), safety margins (cardiac and vascular risk, seizure risk, neurotoxicity signals), and practical deliverability (does it require all-day supervision, does it interact dangerously with common medications, can it be dosed predictably). It also includes an honest plan for bridging the proxy gap: what biomarkers, trial endpoints, and study designs will test whether “reduced psychedelic-like proxy activity” corresponds to “experience-minimized” in humans.

A compound can look excellent on preclinical plasticity and still fail translationally because of toxicity or kinetics. Conversely, a compound can have a clean safety profile but produce weak or inconsistent therapeutic-like effects. “Promise” is not a single axis; it is the intersection of signal, explanation, and feasibility.

For the remainder of the paper, when we say an analogue “shows promise,” we mean it has evidence in all three categories, even if the evidence is early. When it has only one or two, we say so explicitly: “promising preclinical signal,” “promising mechanistic rationale,” or “promising translational direction,” rather than bundling all three into one optimistic word.

3. Case Study: JRT as an LSD Analogue Designed for Reduced Hallucinogenic Potential

3.1 The design premise: minimal structural change, maximal functional shift

JRT is an instructive “stress test” for the popular intuition that psychedelic effects are inseparable from psychedelic scaffolds. The design move was deliberately small: a constitutional re-arrangement that effectively removes the classic indole N–H hydrogen-bonding capability

while leaving the broader ergoline-like geometry and stereochemical “feel” close enough to LSD to plausibly preserve high-affinity binding.

In receptor terms, the premise was not “make LSD weak,” but “make LSD different.” The working hypothesis was that disrupting a specific hydrogen-bond interaction associated with LSD’s ability to stabilize particular 5-HT2A receptor conformations could shift the receptor’s state occupancy away from those most associated with hallucinogenic signaling, while retaining other interactions that support potency and brain-relevant engagement. This is a subtle but important philosophical stance: the target is not simply the receptor, but the receptor-state landscape. JRT is therefore best read as a probe of whether “trip” is a necessary companion of “plasticity,” or whether those two can be pulled apart by precise edits that change how a ligand moves through receptor conformations over time.

3.2 Synthesis overview at the concept level: why de novo matters

A key practical point is that JRT is not “LSD with a tweak you can do late.” Because JRT is a constitutional isomer rather than a simple substitution, it cannot be obtained by straightforward semisynthetic modification of lysergic acid. That matters for two reasons.

First, it forces the chemistry to become part of the scientific claim. A de novo total synthesis is not just a technical flex; it is what makes the test possible. If the only accessible analogues are those reachable by late-stage edits, then the design space is artificially constrained to what synthetic convenience permits, not what mechanistic hypotheses demand.

Second, de novo routes clarify what is actually being optimized. In concept terms, the strategy is modular: build a suitably substituted heteroaromatic fragment and a complementary partially saturated ring fragment, couple them in a controlled way, and then close the remaining architecture while preserving the stereochemical requirements that make LSD-like binding plausible. The reported synthesis also matters as a reproducibility anchor: the structure can be confirmed unambiguously (including stereochemical assignment), and the final compound’s identity does not rest on inference alone. For drug-discovery readers, the big lesson is that “therapeutic without trip” design often depends as much on synthetic seen-it-through-to-the-end feasibility as on receptor theory.

3.3 Cellular signaling profile: what was optimized and what changed

At the cellular pharmacology level, JRT’s profile illustrates a multi-parameter shift rather than a single on/off switch.

One shift is selectivity and polypharmacology. Compared with LSD’s famously broad interaction footprint across multiple receptor families, JRT was characterized as being far more concentrated in serotonergic space, with strong engagement of 5-HT2-class receptors and

comparatively little activity at many non-serotonin CNS targets often implicated in side effects (and in confounds when interpreting behavioral assays). This matters because "reduced hallucinogenic potential" is not only about 5-HT2A—it is also about avoiding additive liabilities that can muddy interpretation or worsen tolerability.

A second shift is efficacy patterning at key serotonin receptors. JRT behaves like a partial agonist at 5-HT2A (and also partial at 5-HT2B), while approaching near-full agonism at 5-HT2C. That combination is mechanistically suggestive: 5-HT2A engagement is still present (consistent with preserving psychoplastogen-like plasticity), but the "gain" on some hallucinogenesis-associated transduction is reduced, while 5-HT2C agonism may contribute to antipsychotic-relevant or dopamine-modulatory behavioral effects in certain paradigms.

A third shift is kinetic and conformational signaling behavior. Rather than relying only on traditional downstream assays (which can be distorted by receptor reserve and pathway amplification), the work leaned on receptor-conformation–proximal readouts and ligand kinetic comparisons. JRT dissociates from 5-HT2A more rapidly than LSD, and its ability to recruit beta-arrestin–linked signaling is markedly reduced compared with LSD, even while it can still drive Gq-linked activation. This combination fits a coherent mechanistic story: changing a specific binding interaction can speed exit from a binding pose that stabilizes certain "long-lived" active conformations and, as a consequence, weaken the subset of signaling behaviors most associated with hallucinogenic-like outputs—without erasing receptor engagement altogether.

Taken together, the cellular profile of JRT is not "LSD turned off." It is "LSD re-tuned": altered contact chemistry, altered kinetics, altered pathway balance, and a different conformational fingerprint—especially when considering differences between human-receptor and mouse-receptor contexts.

3.4 Mouse behavioral proxies: interpreting "reduced hallucinogenic potential"

The paper's behavioral argument hinges on proxies—useful, but intrinsically limited.

The flagship proxy is the mouse head-twitch response (HTR), commonly used because it correlates well with the hallucinogenic potency of many serotonergic psychedelics across species. In that assay, JRT does not produce a robust HTR across tested doses, while LSD does. That by itself supports "reduced hallucinogenic-like activity" in a standard preclinical framework. The more interesting result is antagonist-like behavior in the same proxy: pretreatment with JRT can block LSD-induced HTR. This suggests not merely absence of a hallucinogen-like signal, but a capacity to interfere with the receptor-state or receptor-access pattern that produces it under LSD—consistent with partial agonism plus competitive occupancy and/or conformational competition at 5-HT2A.

Beyond HTR, the work emphasizes not worsening psychosis-relevant proxies. Prepulse inhibition (PPI), a sensorimotor gating measure with translational relevance to psychosis-related phenotypes, is not impaired by JRT; and JRT does not exacerbate pharmacologically induced PPI deficits (for example, those induced by NMDA-antagonist challenges). Additional locomotor paradigms are used to probe dopaminergic and glutamatergic disturbance models: JRT does not itself create the kinds of gross behavioral disturbances that would suggest a problematic psychotomimetic profile, and in at least some contexts it attenuates hyperlocomotion induced by dopaminergic challenge in a sex-dependent manner (a reminder that "reduced risk" claims have to be stress-tested across sex, strain, and model).

The crucial interpretive discipline is this: "reduced hallucinogenic potential" in mice means "reduced signal in validated proxies that track with hallucinogenic compounds," not "proof of no altered perception in humans." The work strengthens the former claim substantially; it cannot, by design, fully establish the latter.

3.5 Antidepressant-like and cognitive-relevant assays: what they suggest, what they do not

JRT's therapeutic promise, in the paper's framing, is not merely "it doesn't twitch." It is "it still behaves like a psychoplastogen."

On the antidepressant-like side, the work reports robust effects in the forced swim test (FST) at post-dose time points consistent with rapid-acting interventions. Interpreting FST responsibly requires restraint—it is a predictive screen with known limitations, not a model of depression—but it remains a standard bridge between mechanistic hypotheses and early behavioral signal.

More compelling, from a translational standpoint, is the use of stress paradigms that produce measurable deficits in reward responsivity and hedonic-like behavior, followed by rescue with JRT. In a sucrose preference framework after chronic corticosterone exposure, JRT rescues stress-induced loss of preference in animals classified as stress-sensitive by baseline-to-post-stress change. This is stronger than a one-off behavioral effect because it anchors the "benefit" to a defined stress-driven deficit.

Even more strategically, the work uses a probabilistic reward task (PRT) adapted for rodents, explicitly motivated by cross-species translational logic: reward learning and response bias deficits are quantifiable in humans and are implicated in depression-related anhedonia. In a chronic stress condition (cold-water stress), JRT rescues reduced response bias in a way that persists for days after a single dose, even with ongoing stress exposure. That persistence—paired with pharmacokinetic evidence that the compound clears relatively quickly—supports the psychoplastogen narrative: brief pharmacological engagement triggers longer-lived circuit-level consequences.

On the cognitive side, reversal learning after unpredictable mild stress (UMS) provides a cognitively legible phenotype: stress impairs flexibility; JRT rescues it without simply boosting baseline discrimination performance. That pattern matters because it argues for "repair under deficit" rather than "nonspecific performance enhancement," and it positions JRT as relevant not only to mood symptoms but to cognitive symptoms often underserved by existing therapeutics.

What these assays do not prove is equally important. They do not prove human antidepressant efficacy, nor do they establish that subjective psychedelic experience is absent in humans. They do, however, establish a credible preclinical case that (a) plasticity-linked and therapeutic-like signals can survive the removal of strong hallucinogen-like proxy signals, and (b) the "benefit" claim is not confined to one assay family.

3.6 The core lesson from JRT: small edits can re-route receptor-state landscapes

JRT's central lesson is architectural: psychedelic pharmacology is not a monolith. A molecule can remain near enough to a psychedelic scaffold to preserve key binding interactions and plasticity-promoting capacity, yet differ enough—through a single, hypothesis-driven structural edit—to alter receptor kinetics, conformational preference, and pathway weighting.

This reframes "therapeutic without trip" as a receptor-state engineering problem. The relevant question becomes: which receptor conformations (and which temporal patterns of occupancy and disengagement) correlate with structural plasticity, and which correlate with perceptual disruption and associated risks? JRT suggests those sets are not perfectly overlapping.

A second lesson is methodological humility: the more "minimal" the chemical change, the more confident we must be that our assays are measuring what we think they measure. JRT looks like a success in mouse proxies and in multiple therapeutic-relevant paradigms, but it also highlights the gap between (i) proxy-defined hallucinogenic potential and (ii) human phenomenology. The right takeaway is neither hype nor dismissal, but a disciplined next-step agenda: human-relevant receptor context, careful translational biomarkers, and clinical designs that do not smuggle certainty into language.

Finally, JRT implies an opportunity for medicinal chemistry strategy: don't only search for new scaffolds—learn to edit known scaffolds with mechanistic intent. The future "without trip" pipeline may be built less by abandoning psychedelic scaffolds and more by mastering the micromoves that redirect them.

4. The Mechanistic Hypothesis Space

The “therapeutic without trip” program lives inside a hypothesis space, not a settled mechanism. The field is trying to separate two bundles of effects that have historically traveled together: (1) acute phenomenology (altered perception, self-processing, salience, meaning) and (2) durable post-acute change (mood, reward learning, cognitive flexibility, relapse liability). The strong version of the hypothesis says the second bundle can be achieved with minimal or no first bundle. The weaker version says the second bundle can be achieved with a smaller, shorter, or qualitatively different first bundle, and that this reduction is clinically valuable even if not absolute.

Mechanistically, the program is often narrated as “optimize 5-HT2A signaling to get plasticity without hallucinations.” That is directionally useful but incomplete. The reality is a web: receptor subtype mix, biased signaling, brain-region specificity, circuit state, developmental history, dosing kinetics, and learning context. This section lays out the main levers the field believes it can pull, and the main ways those levers can fail.

4.1 5-HT2A engagement as a lever, not a guarantee

A core observational anchor of psychedelic science is that many classic psychedelics rely heavily on 5-HT2A engagement for their characteristic acute effects. That makes 5-HT2A a lever: manipulate 5-HT2A interaction and you can often manipulate the “psychedelic-like” signature. But a lever is not a guarantee. Several reasons explain why.

First, 5-HT2A is widely expressed across cortical layers and cell types, and the downstream consequences depend on where and how the receptor is engaged. A ligand that strongly activates 5-HT2A in one microcircuit might promote synaptic remodeling, while engagement in another microcircuit might amplify sensory prediction errors and destabilize perception. Two ligands can share a receptor target and still shift different parts of the brain’s network dynamics.

Second, 5-HT2A effects are strongly state-dependent. The same pharmacological push can land differently depending on baseline arousal, stress hormones, sleep debt, inflammation, prior trauma, ongoing medications, and learned expectations. This matters because “therapeutic” and “psychotomimetic” can be neighboring outcomes separated by context, dose, and vulnerability.

Third, 5-HT2A engagement is often necessary but not sufficient for the desirable end states. Durable improvement likely requires coordinated changes across multiple systems: serotonergic modulation, glutamatergic plasticity machinery, inhibitory interneuron regulation, dopaminergic salience assignment, and network-level reweighting. A compound can hit 5-HT2A cleanly and still fail to produce durable change if the broader cascade does not assemble.

So the mechanistic stance of this paper is conservative: treat 5-HT2A as a privileged entry point into the psychedelic design space, but do not treat it as a magic switch that guarantees either plasticity or safety.

4.2 Biased agonism and downstream pathway selection

If 5-HT2A is the lever, biased agonism is the engineering theory for how to pull it differently. The idea is that receptor activation is not a single event but a choice among downstream pathways. Ligands can stabilize different receptor conformations, which can favor different G-protein couplings, arrestin recruitment patterns, and intracellular signaling cascades. The promise is separation: keep the pathway components that support plasticity while reducing the components that correlate with hallucination-like proxies or psychosis-relevant signatures.

There are three practical subtleties that keep biased agonism from becoming a slogan.

First, “bias” depends on the assay. A ligand can look biased in one cellular context and not in another because receptor density, reveals of receptor reserve, and the expression of signaling partners vary by cell type. A compound “biased” in a recombinant cell line might behave differently in cortical pyramidal neurons or in a mixed network.

Second, bias is not always stable across doses. At low doses a ligand may preferentially engage one pathway; at higher doses it may saturate multiple pathways and collapse back toward a more global activation pattern. The field therefore needs dose-aware bias claims, not categorical ones.

Third, bias interacts with kinetics. A short-lived receptor engagement might produce an intracellular pulse that favors certain cascades, whereas a long residence time might sustain signaling and recruit secondary processes. Even if two ligands have similar pathway preference at a single time point, their time courses can produce different integrated effects.

Despite these caveats, biased agonism remains one of the most plausible mechanistic routes to “therapeutic without trip,” because it gives a structured way to explain how a small structural edit can yield a large functional shift without abandoning the scaffold entirely.

4.3 Plasticity as a multi-node phenomenon: spines, synapses, circuits

“Plasticity” is often spoken of as if it were a single property, but it is a multi-node phenomenon spanning scales.

At the smallest scale, there are structural changes: dendritic spine growth, spine stabilization, and pruning. These changes are not uniformly “good.” Excess growth can be noisy; poorly stabilized growth may be transient; and growth in the wrong circuit can be maladaptive. The relevant question is not “more spines,” but “right spines, right places, right persistence.”

At the synaptic scale, there are functional changes: altered receptor trafficking, changes in excitatory and inhibitory balance, shifts in short-term vs long-term plasticity rules, and changes in how networks encode prediction error. In therapeutic narratives, plasticity is sometimes treated as reopening a developmental window. Mechanistically, it may be better treated as temporarily changing the brain's learning rate and constraint regime.

At the circuit scale, the question becomes integration. Durable symptom relief likely requires that microstructural changes cohere into network-level remapping: altered default network coupling, improved top-down regulation of limbic reactivity, restored reward learning, and improved cognitive flexibility. This is why simple plasticity metrics can mislead: a compound can promote structural plasticity but still fail therapeutically if circuit-level integration is not guided by context, behavior, and learning.

The "therapeutic without trip" hypothesis often assumes that the trip is not needed for plasticity. A more careful version is: the trip may not be needed to trigger plasticity, but some form of guidance may be needed to channel plasticity into useful circuit-level change. That guidance could be psychotherapy, behavioral activation, social support, training tasks, or simply life context. In other words, the experience may be removable, but the learning environment may not be.

4.4 Psychosis-relevant signatures vs therapeutic-relevant signatures

A major safety objective in this design space is to reduce psychosis-relevant risk signals while preserving therapeutic-relevant signals. The field often uses two kinds of signatures as partial guides.

One class involves psychosis-relevant proxies and molecular signatures. These may include disruptions in sensorimotor gating, gene-expression programs associated with dysregulated cortical excitation, and behavioral patterns that resemble overstated salience assignment or disorganized responsiveness under challenge paradigms. These are not psychosis, but they are treated as early warning indicators.

A second class involves therapeutic-relevant signatures: restoration of reward sensitivity under stress, improved cognitive flexibility, reduction of helplessness-like behaviors, and changes in markers associated with synaptic remodeling and stress-circuit normalization.

The mechanistic question is whether these two signature classes are separable. If they are tightly coupled, then reducing one will reduce the other, and "therapeutic without trip" becomes wishful. If they are partially separable, then medicinal chemistry has room to maneuver.

There is a further complication: some pathways that look “therapeutic-relevant” in one context can look “psychosis-relevant” in another. Increased cortical plasticity might help a depressed brain escape rigid negative priors, but in a vulnerable brain it might destabilize already-fragile inference. The same intervention can be reparative or destabilizing depending on baseline constraint and vulnerability.

Therefore, the paper treats these signatures as a balancing act rather than a binary. A promising analogue does not merely “avoid hallucinogenic proxies”; it also shows no meaningful worsening of psychosis-relevant proxies under stress or challenge, while preserving therapeutic-relevant outcomes under deficit conditions. The “under deficit conditions” clause matters: repair under stress is more convincing than enhancement in a healthy baseline.

4.5 Dose, exposure, and kinetics: why timing can mimic “mechanism”

One of the easiest ways to misunderstand this field is to confuse a kinetic effect with a mechanistic effect. Many observed differences between compounds that share a target profile can arise because of differences in how the brain experiences the drug over time.

Peak concentration matters. A rapid spike can push networks into destabilized regimes, producing intense acute phenomenology, while a slower rise can feel calmer even if the receptor engagement is comparable in total.

Duration matters. A long duration can increase both therapeutic opportunity and liability: more time for plasticity, but also more time for anxiety, dysphoria, sleep disruption, and physiological stress. Shorter-acting compounds can be easier to integrate into clinical workflows and can reduce acute burden, but they may not deliver the same integrated downstream effect if duration is part of what recruits the plasticity cascade.

Exposure matters. Total area-under-the-curve can be similar across compounds while subjective intensity differs, and vice versa. A compound that binds tightly and dissociates slowly can behave like a sustained “pressure” on the receptor landscape, potentially increasing the probability of entering certain receptor states or recruiting secondary signaling and transcriptional programs.

Timing also interacts with learning. If a compound creates a plasticity window, the timing of therapeutic input relative to that window may matter more than the identity of the compound. A therapy delivered at the wrong time can make an effective compound look ineffective, or make a marginal compound look effective by coincidence.

For “therapeutic without trip,” kinetics are not merely an optimization detail; they are part of the hypothesis. A compound might reduce hallucinogenic-like effects not because it lacks the relevant signaling capacity, but because it avoids the rapid, high-amplitude perturbations that

produce strong subjective and proxy signatures. Conversely, a compound might appear “clean” in proxies at one exposure profile and not at another.

The practical conclusion is that mechanism claims must be dose- and time-aware. When two analogues differ, the first interpretive question should be: is the difference explained by pathway selection at the receptor, or by the brain’s time-resolved exposure to the ligand? Often, the honest answer is “both,” and the design space is precisely the interaction between these two.

5. The Assay Stack: How the Field Decides a Compound Is “Less Psychedelic”

The phrase “less psychedelic” sounds like an experiential statement, but in the scientific workflow it is usually a conclusion drawn from a stack of assays—some mechanistic, some behavioral, some translationally motivated, some simply conventional. The assay stack functions like a set of filters. A candidate compound is advanced if it clears the early filters (binds the right receptors, has a plausible signaling profile), then clears mid-level filters (shows desirable plasticity and therapeutic-like behavior), and then does not trigger red flags (hallucinogen-like proxies, safety liabilities). The stack is not standardized across labs, but it follows recognizable patterns.

The central discipline is to remember what the stack is doing: it is not “proving non-hallucinogenic.” It is triaging candidates based on correlates and proxies. When the stack is used responsibly, it produces statements like “reduced hallucinogen-like proxy activity at exposures that retain plasticity and therapeutic-like effects.” When it is used irresponsibly, it is translated into public language as “non-hallucinogenic psychedelic,” which is a much stronger claim than the stack can support.

5.1 Behavioral proxies (for example, head-twitch response): strengths and limits

Behavioral proxies are often the first public-facing evidence because they are legible: you can show a plot of “number of head twitches” and it looks like a behavioral signature of “trip intensity.” The most famous is the mouse head-twitch response (HTR), widely used because it correlates with 5-HT2A activation for many serotonergic psychedelics and tracks the potency ranking of classic psychedelics surprisingly well.

Strengths of behavioral proxies include:

- Practicality and throughput. HTR can be measured relatively quickly across multiple doses and compounds, making it useful for screening.

- Pharmacological anchoring. Antagonists that block 5-HT2A often block HTR, supporting the idea that the proxy is mechanistically connected to the receptor that matters.
- Comparative utility. When used consistently, HTR provides a common yardstick for comparing analogues.

However, the limits are not minor—they define what can and cannot be claimed.

First, proxies are not subjective experience. An animal is not reporting perception; it is producing a stereotyped motor action that correlates with certain receptor-driven states. Absence of HTR is not proof of absence of altered consciousness in humans.

Second, HTR is sensitive to dose, strain, sex, environment, and methodological details. A compound may appear "clean" in one context and not in another. Small shifts in experimental conditions can change baseline movement, detection sensitivity, and scoring.

Third, HTR is most informative for serotonergic psychedelics and less informative as you move into mechanistically mixed compounds. If an analogue engages multiple receptor families, the HTR signal can be masked or modulated by sedative, stimulant, or motor-disrupting effects that confound interpretation. A compound might reduce HTR simply because it reduces movement, not because it reduces psychedelic-like signaling.

Fourth, HTR focuses attention on one receptor axis and one behavioral readout. Psychedelic phenomenology is a high-dimensional state. A compound might reduce the HTR while still producing strong changes in affect, salience, time perception, or dissociation-like effects. If the field equates "no HTR" with "no trip," it is compressing complexity into a single number.

Therefore, in this paper, behavioral proxies are treated as necessary screening tools and useful correlates, but never as definitive statements about human experience. The correct inference is: "in this proxy, the compound behaves unlike classic serotonergic psychedelics at these doses," and that inference must remain conditional and dose-specific.

5.2 Gene-expression and circuit readouts: what they add beyond behavior

Gene-expression and circuit-level readouts are the field's attempt to move from "what the mouse did" to "what the brain is doing." They add three kinds of value.

First, they can distinguish plasticity programs from stress-toxicity programs. A compound that causes acute behavioral changes might do so by sedation, agitation, or stress, while a compound that triggers plasticity-related transcriptional cascades may show a different pattern of immediate early genes and downstream synaptic remodeling markers. These signatures can help separate "behavioral artifact" from "biological remodeling."

Second, they can reveal divergence between compounds that look similar on behavior. Two analogues might both reduce HTR, but one might strongly activate cortical transcriptional programs linked to psychedelic-like states while the other does not. Conversely, two compounds might both show antidepressant-like behavior, but one might do so without activating psychosis-relevant gene-expression programs under challenge conditions. This matters because "less psychedelic" is sometimes less about behavior and more about which brain-wide programs are being recruited.

Third, circuit readouts can connect to translational biomarkers. Electrophysiological signatures, neural oscillations, and functional connectivity changes can be compared across species more plausibly than subjective reports alone. Even if the mapping remains imperfect, circuit-level measures provide a bridge.

The limitation is that these readouts can also become overinterpreted. Gene-expression changes are real but not always legible; a signature can be "different" without being "better." Circuit perturbations can be interpreted as therapeutic normalization or as destabilization depending on theory. A responsible stack uses these readouts as triangulation: they support or challenge the behavioral story, they do not replace it with a new kind of certainty.

5.3 Receptor profiling and functional selectivity: common pitfalls

Receptor profiling and functional selectivity assays are the molecular foundation of the stack, and they are also where misunderstandings begin.

A typical workflow includes binding affinity panels, functional assays at key receptor subtypes, and sometimes more conformation-proximal methods. These data are used to argue: "The compound targets 5-HT2A in a way consistent with plasticity, but does not engage other targets likely to cause side effects or confounds."

Common pitfalls include:

- Confusing affinity with effect. High affinity does not imply high efficacy, and low efficacy does not imply low impact if the receptor is highly expressed or the downstream amplification is strong.
- Ignoring receptor reserve and assay context. A partial agonist can look like a full agonist in a high-reserve system. Bias claims can flip depending on receptor density and pathway coupling.
- Overstating "bias" from a single assay. Bias is not a property you can declare once; it must be demonstrated across relevant contexts and ideally in neuron-like systems or ex vivo tissue where the signaling environment is closer to brain reality.

- Failing to control for kinetics. A ligand with long residence time can produce sustained signaling and secondary effects that look like “different pathways” when the difference is actually time-integrated occupancy.
- Overlooking species differences. A compound can behave differently at mouse and human 5-HT2A receptors. If the goal is human translation, human receptor data must be treated as primary.

The point of these pitfalls is not to dismiss receptor profiling; it is to keep it honest. Medicinal chemistry needs mechanistic handles. But the handles must be attached to the right object: not just receptor identity, but receptor-state dynamics in relevant biological contexts.

5.4 Safety pharmacology: cardiotoxicity, vasoconstriction, seizure risk, abuse liability

If “therapeutic without trip” is an attempt to remove one kind of liability (intense altered experience), it does not automatically remove others. In fact, some liabilities may become more important if compounds are intended for broader, more routine use.

Key safety domains include:

- Cardiotoxicity and arrhythmia risk. Compounds that interact with cardiac ion channels or provoke dangerous rhythm changes can fail regardless of psychiatric promise.
- Vasoconstriction and vascular risk. Certain serotonergic receptor interactions can produce vasoconstriction and blood pressure effects. If a compound is to be used widely, even modest vascular liabilities become important.
- Seizure risk. A compound that alters excitation–inhibition balance in cortical networks can have proconvulsant potential under some conditions.
- 5-HT2B and valvular concerns. Because 5-HT2B agonism has historically been associated with cardiac valvulopathy risk for certain drugs, its engagement is often examined carefully. A compound can look therapeutically exciting and still be clinically constrained if it engages 5-HT2B in a problematic way over chronic exposure.
- Abuse liability and behavioral reinforcement. A “no-trip” profile does not guarantee low abuse potential; a compound could be reinforcing for other reasons (stimulation, anxiolysis, mood elevation). Conversely, some classic psychedelics have low compulsive use patterns, but analogues can shift that profile unpredictably.

A major strategic challenge is that the desired therapeutic model matters for safety interpretation. If a compound is intended for rare, supervised administration, some acute risks can be managed clinically. If a compound is intended for more frequent or outpatient use, the

tolerability and safety bar becomes much higher. “Therapeutic without trip” often implicitly aims at scalability, and scalability raises safety expectations.

5.5 A proposed evidence ladder: from receptor to rodent to human

Because the field’s language often outruns its evidence, it helps to formalize a ladder of claims—what you are allowed to say at each rung, and what you are not.

Rung 1: Molecular engagement.
You can say: the compound binds and functionally engages specific receptors with defined potency and efficacy in defined assay systems.
You cannot say: it will be therapeutic, or it will be non-hallucinogenic in humans.

Rung 2: Signaling and bias profile.
You can say: in these assays and contexts, the compound shows pathway preference, reduced arrestin recruitment, altered kinetics, or other mechanistic differences relative to a benchmark psychedelic.
You cannot say: the compound has separated plasticity from perception in humans.

Rung 3: Plasticity markers.
You can say: the compound increases structural and/or functional plasticity markers in neurons or brain tissue and shows time courses consistent with remodeling.
You cannot say: this plasticity is automatically beneficial or clinically relevant.

Rung 4: Behavioral efficacy under deficit.
You can say: the compound produces antidepressant-like, anti-addictive, or cognitive flexibility effects in preclinical models, preferably under stress or deficit conditions, with appropriate controls.
You cannot say: it treats depression or addiction in humans.

Rung 5: Reduced psychedelic-like proxy activity.
You can say: the compound produces reduced signals on established proxies (for example, reduced HTR), and does not worsen psychosis-relevant proxies under challenge.
You cannot say: it is non-hallucinogenic in humans.

Rung 6: Early human pharmacology and phenomenology.
You can say: in controlled human studies, the compound produces minimal subjective psychedelic effects at exposures associated with biomarker or symptom changes, and shows tolerability.
Only here can you responsibly approach the phrase “experience-minimized,” and even then you must specify dose, population, and measurement.

Rung 7: Clinical efficacy and durability.
You can say: controlled clinical trials demonstrate symptom improvement with defined dosing and acceptable safety.
Only here can the field claim “therapeutic” in the ordinary medical sense.

The ladder clarifies why social media often misleads: it jumps from rung 5 to rung 7. The correct narrative is slower, but it is also stronger. If the field wants durable trust, it must learn to speak in rung-aware language.

6. Other Analogues Showing Promise

The JRT case study is powerful because it compresses the “therapeutic without trip” design idea into a single vivid move: a small structural edit appears to reshape downstream signaling and animal proxy signatures while preserving therapeutic-relevant effects in multiple paradigms. But JRT is not alone. A broader set of candidates—some from academic labs, some from industry pipelines—illustrate that there may be multiple routes to the same objective: preserve plasticity and durable behavioral change while reducing hallucinogen-like proxy signals and practical burdens of administration.

This section is not an exhaustive catalog. It is a curated set of “representative exemplars” chosen because they differ in scaffold family, discovery strategy, and intended clinical niche. The goal is to help the reader see the design space as a landscape rather than a single storyline.

6.1 Tabernanthalog (TBG): an ibogaine-inspired neuroplastogen

Tabernanthalog, often abbreviated TBG, is best understood as a flagship example of “scaffold rescue.” The parent compound, ibogaine, has long been associated with anti-addictive effects and psychologically intense, often destabilizing experiences. It also carries serious safety concerns. The analogue logic here is not merely “make it less hallucinogenic,” but “make it clinically survivable.” That means reducing toxicity liabilities, simplifying pharmacology where possible, and keeping whatever aspect of the parent’s action seems tied to durable therapeutic change.

TBG’s appeal in the “therapeutic without trip” narrative comes from the claim that it can promote plasticity and produce behavioral benefits in addiction- and mood-relevant paradigms while lacking strong hallucinogen-like proxy signatures in the standard rodent assay language. In the best interpretation, TBG suggests that the therapeutic-like effects of a historically “visionary” compound may not require the visionary experience—at least not in the stark form the parent produces.

Mechanistically, TBG also reminds the field of something easy to forget when focusing on 5-HT2A: some of the most interesting anti-addiction signals in nature-derived compounds involve multi-receptor, multi-circuit effects that are not cleanly reducible to a single receptor story. That is both an opportunity and a risk. The opportunity is that non-5-HT2A pathways might deliver plasticity and behavioral change with less perceptual disruption. The risk is interpretability: if a compound has broad action, you can win behavior and lose mechanistic clarity, making it harder to engineer improved second-generation candidates.

Clinically, TBG is often discussed in the addiction lane, but its relevance to mood and stress circuitry is part of its broader appeal. It represents a design philosophy that could be summarized as: build a compound that produces a plasticity window and a behavioral reset without requiring a psychologically intense event as the delivery vehicle. Whether this can be maintained in humans across diverse histories and vulnerabilities remains an empirical question, but as a program it has helped define what "non-hallucinogenic psychoplastogen" is supposed to mean in practice.

6.2 AAZ-A-154 and related "psychLight"-screened candidates

AAZ-A-154 and the "psychLight"-screened class represent a different philosophy: not "rescue a dangerous parent," but "engineer the discovery process itself." Here, the key innovation is methodological. Instead of beginning with a known psychedelic and trying to sand down its rough edges, the approach begins with a receptor-linked biosensor and uses it to search chemical space for ligands that trigger desired signaling signatures without triggering hallucinogen-like proxy signatures.

Conceptually, this is a move from scaffold-first to signal-first discovery.

The promise of the psychLight approach is that it gives the field a new kind of filter at the earliest stage: a way to identify compounds that engage 5-HT2A (and related serotonergic systems) in a more selective and potentially safer signaling mode. The compounds that emerge from this pipeline are often described as "non-hallucinogenic" in the preclinical proxy sense, while still producing plasticity-related cellular changes and durable behavioral effects after a single administration in rodent paradigms.

Two features matter here.

First, the platform suggests that "psychedelic-like" is not a binary but a spectrum of receptor-state and downstream signaling patterns. If your discovery tool can sense those patterns, you can aim medicinal chemistry at a more precise target than "bind 5-HT2A."

Second, the platform hints that durability may not require prolonged psychedelic exposure. Some of the reported candidates are notable precisely because they are framed as producing

long-lasting behavioral changes after transient receptor engagement. This pattern reinforces the psychoplastogen thesis: the drug's acute presence is not the same thing as the duration of its benefit.

The deeper value of the psychLight class for this paper is not the specific molecule name. It is the existence proof that discovery can be retooled around mechanistic aspiration: select for plasticity-linked signatures and select against hallucinogen-like proxies. Even if specific candidates rise and fall with replication and translation, the methodological shift is likely to persist.

6.3 2-bromo-LSD: an ergoline line with recurring clinical interest

2-bromo-LSD is an unusual "bridge object" in this field: it is close enough to LSD structurally to sit inside the lysergamide family, but it has long been described as lacking the classic psychedelic subjective effects in humans at doses explored in certain clinical contexts. It re-enters the modern conversation repeatedly because it offers something the field craves: an ergoline that appears to separate therapeutic-like utility from the psychedelic experience, at least in specific indications.

Historically, it has been associated with headache disorders, especially cluster headache, where small uncontrolled series and clinical interest have kept it alive in the scientific imagination. More recently, it has been used in preclinical work as a comparator to help answer a conceptual question: can an LSD-like scaffold deliver durable circuit benefits without producing a psychedelic-like behavioral proxy signature?

Even if one brackets the clinical evidence as preliminary or context-specific, 2-bromo-LSD matters as a design hint. It suggests that the ergoline scaffold does not inevitably entail psychedelic phenomenology in the way a casual reader might assume. That alone is valuable, because it helps correct an overly simplistic narrative: "LSD-like structure equals LSD-like trip."

It also matters because it is a sober reminder about translational messaging. If a compound has some human history of "non-psychedelic" characterization, the temptation is to treat that as settled truth and move on. But the disciplined approach is to treat it as a context-bound statement: dose, population, measurement methods, and clinical setting all matter. The right reading is not "this is the answer," but "this is a rare anchor point that can calibrate how we interpret proxy-based claims for newer compounds."

6.4 Lisuride as a comparator and the controversy over "non-psychedelic" labels

Lisuride is frequently invoked as a comparator because it occupies a psychologically provocative niche: it is an ergoline with serotonergic activity that has not been widely associated with classic psychedelic experiences in routine clinical use. This has led many authors to treat it as a kind of

“non-psychedelic 5-HT2A agonist” reference point. If you can understand why lisuride is not a classic psychedelic in practice, you might learn what features of receptor engagement and systems pharmacology actually matter for psychedelic phenomenology.

But lisuride is also a cautionary tale about labels.

First, lisuride’s pharmacology is not “cleanly serotonergic.” It has dopaminergic and other receptor actions that can complicate interpretation, including side effects that limit how it is experienced and tolerated. That matters because a drug can fail to produce a “psychedelic-like” subjective state for reasons that have nothing to do with the absence of psychedelic mechanisms. Nausea, dysphoria, sedation, or other adverse effects can flatten or mask phenomenology in ways that make simplistic inferences unreliable.

Second, the category “non-psychedelic” can become a rhetorical shortcut that prevents careful dose-aware discussion. A compound may be “non-psychedelic” at clinically typical doses but could produce altered states at higher exposures or in susceptible populations, or under particular contexts. Conversely, it may produce subtle but real cognitive and affective shifts that are simply not culturally coded as “psychedelic.”

Third, lisuride underscores a core mechanistic point: 5-HT2A engagement alone does not guarantee a psychedelic experience. That claim, while often stated, is easy to underestimate. Many readers implicitly assume that if you activate 5-HT2A, the trip is unavoidable. Lisuride is one of the reasons the field is increasingly comfortable describing 5-HT2A as a lever rather than a guarantee.

In the context of this paper, lisuride’s main role is conceptual calibration. It helps prevent the field from treating “5-HT2A agonist” and “psychedelic” as synonyms, and it highlights the importance of kinetics, pathway weighting, and polypharmacology in shaping experience.

6.5 Zalsupindole and modern industrial pipelines for non-hallucinogenic neuroplastogens

Zalsupindole represents a newer industry-facing archetype: compounds explicitly positioned as non-hallucinogenic neuroplastogens, designed for scalable psychiatric use. The core idea is similar to JRT’s: preserve plasticity and antidepressant-like benefits while minimizing hallucinogen-like proxy activity and other liabilities that would require intensive supervision.

What distinguishes the modern pipeline approach is less the philosophical aim and more the engineering infrastructure around it:

- Structure-guided design and iteration that treats receptor-state engagement as a tunable property rather than a mystery.

- Broad profiling early to avoid surprises: receptor panels, safety pharmacology screens, and pharmacokinetic studies that are integrated into candidate selection rather than postponed.
- A translational posture that emphasizes measurable endpoints beyond subjective experience: biomarkers, task-based measures, and operational definitions of "experience-minimized."

Industrial candidates also highlight a practical truth: if the goal is wide access, the tolerability bar rises. Classic psychedelic-assisted therapy can manage acute adverse experiences with supervision and preparation. A scalable outpatient-style medicine needs to be more predictable and less demanding on infrastructure. This is one reason industry pipelines are so attentive to minimizing acute phenomenology, shortening or controlling duration, and reducing physiological stress signatures.

The larger point is that "non-hallucinogenic neuroplastogen" is becoming a product category, not only a scientific curiosity. That category will live or die on human data, but the pipeline itself changes what kinds of molecules get built and tested. It increases the probability that multiple candidates will reach early clinical evaluation, which in turn will test whether proxy-defined "less psychedelic" actually maps to human experience and therapeutic durability.

6.6 Why these candidates matter as a set: convergences and divergences

Taken together, these exemplars clarify the shape of the design space.

Their convergences are instructive:

- All treat durable change as the goal and acute phenomenology as optional, modulated, or even unwanted depending on patient and delivery context.
- All rely on a multi-layer assay stack: receptor engagement and signaling, plasticity measures, behavioral efficacy under deficit, and hallucinogen-like proxy suppression.
- All implicitly assume that therapeutic value can be extracted from psychedelic biology without requiring the full psychedelic experience. Even when the programs differ in rhetoric, the design direction is shared.

Their divergences are equally important:

- Different scaffold families may imply different risk profiles and different clinical niches. Ergoline-based candidates often intersect with headache and serotonergic precision questions; ibogaine-derived candidates often intersect with addiction and broader polypharmacology; biosensor-screened candidates often intersect with discovery methodology and signal-first engineering.

- Mechanistic narratives differ. Some programs lean heavily on biased agonism and pathway selection; others lean on kinetics and receptor residence time; others lean on network-level restoration and stress-circuit repair as the core.
- The “trip” is treated differently. In some philosophies it is an undesirable side effect to be eliminated; in others it is something to be minimized for scalability while remaining available in other therapeutic modalities.

As a set, these candidates also teach a discipline of interpretation. The field is drifting toward a more careful vocabulary: not “non-hallucinogenic psychedelic” as a final claim, but “reduced hallucinogen-like proxy activity with preserved plasticity and therapeutic-like outcomes,” followed by the crucial next step: human testing that measures both subjective experience and durable benefit.

The strategic message is that the field should not bet everything on a single mechanism story. Multiple routes may exist from molecule to plasticity window to durable change. The next phase of progress will come from identifying which routes are real, which are fragile, and which generalize across populations. The analogues discussed here matter because they provide multiple, structurally distinct probes of the same question: can we design medicines that open the brain to change without requiring the mind to pass through an overwhelming altered state?

7. Design Principles Emerging Across Scaffolds

As soon as more than one “therapeutic without trip” candidate exists, the field stops being a story about a single clever molecule and becomes a story about principles. Those principles are not yet laws. They are heuristics—pattern-recognition under uncertainty—built from a growing set of analogues that seem to preserve plasticity-linked outcomes while reducing hallucinogen-like proxy signatures and operational burdens. The aim of this section is to surface the recurring design moves that appear across scaffold families, and to translate them into a practical checklist that a medicinal chemistry reader can actually use.

A key theme runs through all of these principles: “separation” is rarely absolute. The more honest goal is controlled decoupling. The field is trying to reduce the probability and intensity of destabilizing perceptual and self-processing alterations at exposures that still trigger durable plasticity and therapeutic-like behavioral change. That is a tunable design objective, not a binary switch.

7.1 Minimize perceptual liability while preserving plasticity: what "separation" looks like

In practice, separation has three meanings, which are often conflated.

First, separation as proxy decoupling. This is the preclinical version: the compound produces robust plasticity markers and therapeutic-like behavioral effects while showing minimal hallucinogen-like proxy activity (for example, low HTR). This is the earliest measurable form of separation and is the most common claim in the literature.

Second, separation as phenomenology reduction. This is the human version: the compound produces minimal subjective psychedelic effects, or produces effects that are qualitatively different—less perceptual distortion, less ego-dissolution, less destabilizing salience—while still producing clinically meaningful benefit. This is the claim many readers assume, but it is usually not yet supported for early candidates.

Third, separation as clinical workflow simplification. This is the health-system version: the compound can be administered without all-day supervision and specialized psychotherapeutic containment, because acute phenomenology and acute risk are sufficiently reduced. This is a practical end point independent of whether some subtle subjective effects remain.

What does separation "look like" when it succeeds?

- The compound retains a plasticity window (structural remodeling markers, circuit restoration signatures).
- Therapeutic-like behavioral benefits persist beyond the presence of drug in the system, supporting the idea of state change rather than symptom masking.
- Hallucinogen-like proxies are reduced relative to benchmark psychedelics at matched exposures.
- Psychosis-relevant proxies are not worsened under challenge paradigms.
- Acute tolerability is improved: lower anxiety-like responses, less dysphoria, fewer autonomic spikes, less sleep disruption.
- The compound has a safety and pharmacokinetic profile consistent with its intended delivery model (single supervised dose vs repeated outpatient use).

The last point is often ignored. A compound can be "less psychedelic" and still be clinically unworkable if it is too long-acting, too unpredictable, or too physiologically stressful. So separation has to be defined against a delivery model, not against ideology.

7.2 Structural strategies: atom swaps, ring edits, substituent tuning, stereochemistry

Across scaffold families, several structural moves recur because they are the most direct way to alter receptor-state occupancy, binding kinetics, and signaling preference without discarding the potency that makes the scaffold interesting.

Atom swaps and heteroatom re-patterning.
The JRT-style move—transposing or swapping atoms in a way that changes hydrogen-bond donors/acceptors—has a disproportionate effect because hydrogen bonding often anchors specific receptor conformations. Removing an indole N–H, shifting a heteroatom position, or changing the directionality of a hydrogen-bond network can change which receptor microstates are stabilized. Importantly, this is not "making the molecule weaker"; it is changing the geometry of information transfer from ligand to receptor.

Ring edits and rigidification/derigidification.
Small changes in ring size, ring fusion, and conformational flexibility can change the ligand's ability to adopt binding poses that favor certain active states. A more rigid ligand may lock the receptor into a specific conformation; a more flexible ligand may sample multiple conformations, sometimes reducing sustained stabilization of any one "psychedelic-favoring" state. Ring edits also affect metabolic stability and brain penetration.

Substituent tuning and electronics.
Substituents change more than affinity. They change electronics, polarity, and the balance between lipophilicity and solubility. They can also influence how deeply the ligand sits in the binding pocket and how it engages extracellular loops that may be important for residence time and arrestin recruitment. The medicinal chemistry pattern is often: keep the "recognition elements" that preserve potency, but tune substituents to weaken or reshape interactions believed to correlate with hallucinogen-like proxy signaling.

Stereochemistry as a functional dial.
Many scaffolds are stereochemically rich. Small changes in stereochemistry can flip functional outcomes because receptor binding is stereoselective and conformation-sensitive. The stereochemical point matters especially in "separation" programs because one enantiomer may preserve desirable signaling while the other adds liabilities. When stereochemistry is treated as a dial rather than a nuisance, it becomes one of the cleanest levers for decoupling effects without adding new atoms or new functional groups.

A practical interpretation of these structural strategies is this: the field is learning to treat psychedelic scaffolds as platforms, not idols. The goal is not reverence for the parent molecule; it is disciplined editing of the parent molecule.

7.3 Signaling strategies: partial agonism, functional selectivity, pathway bias

Structural edits are often justified in terms of signaling strategies—what the ligand does after it binds.

Partial agonism as risk reduction.
A partial agonist can occupy the receptor and produce sufficient signaling to trigger plasticity-related cascades without driving the receptor system into the high-amplitude regimes that correlate with strong psychedelic proxy signatures. Partial agonism can also confer antagonist-like properties in the presence of a full agonist, which is one reason some analogues can block hallucinogen-like proxies induced by classic psychedelics.

Functional selectivity and pathway weighting.
The strongest mechanistic story for separation is not “less signaling,” but “different signaling.” Ligands that reduce beta-arrestin recruitment relative to Gq activation, or that shift downstream phosphorylation patterns, are often framed as emphasizing plasticity-relevant cascades while de-emphasizing phenomenology-linked cascades. Whether that mapping is fully correct remains under test, but as a design principle it is coherent: choose receptor-state stabilization that changes the downstream mixture.

Pathway bias must be time-aware and context-aware.
Bias is not a static property. It can vary with dose, receptor density, cell type, and time course. A compound that appears “biased” in a recombinant cell line may behave differently in cortical tissue. For separation programs, bias claims should be treated as provisional until they are demonstrated in neuron-relevant systems and under exposure profiles consistent with in vivo dosing.

The practical consequence is that signaling strategies must be specified with constraints: which pathway, in which system, at which time point, at which exposure. Without those constraints, “biased agonism” becomes a banner rather than a mechanism.

7.4 Systems strategies: kinetics, brain penetration, metabolite control

Even perfect receptor engineering can be undermined by system-level behavior. Many “trip” liabilities are not strictly receptor identity problems; they are exposure problems.

Kinetics as phenomenology shaping.
Fast onset often correlates with intensity. A rapid spike can produce a destabilizing acute perturbation, while a slower rise can reduce anxiety and perceptual disruption even if total receptor engagement is similar. Duration also matters: long duration increases both therapeutic opportunity and acute burden. Many separation programs implicitly aim for “shorter, cleaner, predictable” exposure profiles.

Brain penetration and regional exposure.
Two compounds with similar plasma levels can produce different brain effects if their brain penetration differs, if they partition differently into membranes, or if transporters and metabolism shape their regional distribution. A “less psychedelic” profile could arise because a compound avoids peak cortical exposure while still engaging subcortical or other targets relevant to mood and reward. This is not guaranteed, but it is a plausible route.

Metabolite control.
Metabolites can be active, and active metabolites can defeat design intentions. A parent compound designed to be experience-minimized can metabolize into a compound that is not. Conversely, a parent compound can act as a prodrug that smooths exposure and reduces acute intensity. In either case, metabolite profiling is not optional if the field wants clean language.

Delivery model alignment.
A compound intended for rare supervised dosing can tolerate certain kinetic features that would be unacceptable for outpatient use. Separation programs often want outpatient viability, so they must control half-life, inter-individual variability, drug–drug interactions, and food effects. In the end, “less psychedelic” often means “less variable,” because unpredictability itself is a liability.

Systems strategies therefore act like a second layer of design: even if receptor-state engineering is successful, the compound must deliver that engineering into the brain in a controlled way.

7.5 A practical medicinal chemistry checklist for “safer psychoplastogens”

This checklist is deliberately practical. It is meant to be a discipline against hype: a way to ask “what exactly have we shown?” and “what did we not show?” for any new analogue claim.

Structural and chemical identity

- Confirm unambiguous structure and stereochemistry.
- Demonstrate synthetic reproducibility and purity (including isomer control).
- Characterize stability and major metabolic pathways early.

Target engagement and selectivity

- Measure functional activity at human 5-HT2A and key serotonergic subtypes (including 5-HT2B and 5-HT2C).
- Profile off-target liabilities likely to confound behavior or safety (cardiac ion channels, major CNS receptors).
- Avoid claiming “selective” without showing the panel and the exposure context.

Signaling characterization

- Quantify efficacy and bias using multiple readouts, not one.
- Test bias in neuron-relevant systems when possible, not only recombinant cell lines.
- Include kinetic measurements (residence time, association/dissociation) when the scaffold suggests long-lived binding.

Plasticity evidence

- Show structural plasticity markers (spines, synaptic proteins) with time course and persistence.
- Demonstrate that plasticity occurs at exposures that do not trigger strong proxy signals.
- Distinguish plasticity from generalized stress responses.

Behavioral efficacy under deficit

- Prefer models where a deficit is induced (stress, learned helplessness-like paradigms, reward learning deficits, relapse models).
- Look for persistence beyond drug clearance to support state change rather than masking.
- Confirm that benefits are not explainable by sedation, stimulation, or motor artifacts.

Hallucinogen-like proxy profile

- Test at doses matched for brain exposure, not only matched by mg/kg.
- Report multiple proxy readouts when possible (HTR plus at least one additional measure).
- Avoid translating “no HTR” into “no trip” without human data.

Psychosis-relevant risk checks

- Assess sensorimotor gating and challenge paradigms where appropriate.
- Report sex-dependent and strain-dependent effects explicitly, not as afterthoughts.
- Treat “no worsening” as a minimum requirement, not as a victory lap.

Safety pharmacology and tolerability

- Evaluate cardiovascular and vascular risk signals, including 5-HT2B concerns in relation to dosing frequency and exposure duration.

- Screen seizure liability and CNS adverse effect patterns.
- Address abuse liability explicitly, especially if the compound will be outpatient.

PK/PD and translational planning

- Establish brain exposure, half-life, variability, and metabolite activity.
- Propose a biomarker bridge plan for early humans: what will be measured besides subjective reports?
- Match the compound's profile to a realistic clinical delivery model.

Claims discipline

- Use rung-aware language: "reduced hallucinogen-like proxy activity" is not "non-hallucinogenic in humans."
- State what was measured, in what species, at what dose, and for how long.
- Make "unknowns" explicit, especially regarding human phenomenology.

If a candidate clears most of this checklist, it deserves serious attention. If it clears only a few items but is discussed as a finished product, the field is likely watching a narrative inflate faster than evidence. The role of these design principles is to keep innovation alive while keeping language honest—so that when a true "safer psychoplastogen" arrives, it will be recognized as such for the right reasons.

8. Translational Strategy: From Animal Proxies to Human Claims

Translation is the hard part of this entire enterprise. In preclinical work, "less psychedelic" is mostly a statement about proxies. In the clinic, it becomes a statement about subjective experience, safety, and durable benefit under real-world constraints. The bridge between these domains is not automatic; it must be engineered. The field needs a translational strategy that is explicit about what can be carried across species (mechanistic signatures, circuit perturbations, plasticity markers), what cannot (subjective phenomenology), and how to measure outcomes without smuggling assumptions into language.

A mature translational posture also recognizes a key asymmetry: false positives are cheap in rodents and expensive in humans. A compound can look elegant in proxy space and fail immediately when people report unexpected phenomenology or when benefits do not materialize. The point of translation is therefore not to prove a story; it is to design tests that can falsify it quickly and safely.

8.1 Biomarkers that might bridge species (EEG, imaging, peripheral readouts)

A bridge biomarker has a specific job: it should be measurable in both animals and humans (or at least have homologous forms), it should plausibly relate to the hypothesized mechanism (plasticity windows, network reweighting, stress-circuit normalization), and it should be interpretable without relying on the patient's narrative of the session. That last clause is central for "therapeutic without trip," because the program is explicitly trying to avoid treating the subjective experience as the mechanism of action.

EEG and related electrophysiology.
EEG is attractive because it is accessible, relatively inexpensive, and time-resolved. It can capture acute changes (during drug exposure) and post-acute changes (hours to days later). For classic psychedelics, EEG often shows characteristic changes in oscillatory power and signal complexity; the exact patterns can vary, but the general idea is that network dynamics become more "entropic" and less dominated by rigid rhythms. For experience-minimized analogues, the hypothesis is nuanced: the compound may produce a smaller acute EEG perturbation while still producing a measurable post-acute shift consistent with altered plasticity or network reweighting.

In translation terms, the key is to predefine what "bridging" means. A bridging EEG signature might be a reproducible post-dose change in specific band power relationships, coherence metrics, or complexity measures that correlates with behavioral improvement or with other biomarkers of plasticity. The field must be cautious here: high-dimensional EEG can generate patterns that look meaningful by chance. The biomarker must be pre-registered, and its interpretation must be conservative.

Imaging: fMRI, PET, and structural measures.
Imaging can serve different translational roles. fMRI can reveal changes in functional connectivity, especially in networks implicated in rumination, salience assignment, and cognitive control. PET can quantify receptor occupancy and, in some contexts, downstream neurochemical changes. Structural imaging can capture longer-term changes, though those are often subtle and confounded by time and environment.

For "therapeutic without trip," imaging is attractive because it can decouple subjective reports from network-level changes. But it is also risky: imaging signatures are easy to overinterpret. A useful translational imaging approach is hypothesis-driven and narrow. Examples include:

- Demonstrate dose-dependent receptor occupancy that matches the exposure range where preclinical plasticity markers are observed.

- Identify a post-acute network normalization signature (for example, reduced hyperconnectivity of certain networks under stress) that persists beyond acute exposure.
- Correlate imaging changes with task performance or symptom improvement, not only with self-report of subjective experience.

Peripheral readouts and blood-based biomarkers.
Peripheral measures are tempting because they scale. Candidates include markers related to synaptic plasticity pathways, inflammatory tone, stress-hormone regulation, and neurotrophic signaling. The field often discusses neurotrophin-related markers, but peripheral levels are not direct windows into the brain. The real value of peripheral biomarkers is not that they "prove plasticity," but that they can support a coherent multi-modal story when paired with EEG, imaging, and behavioral outcomes.

A pragmatic translational strategy might use peripheral biomarkers as "state indicators" rather than mechanistic proof: signs that the system has entered a different biological regime after dosing, which can then be related to symptom and function changes.

Task-based behavioral biomarkers.
If the goal is to avoid "trip endpoints," the most valuable human measures may be task-based. Reward learning tasks, cognitive flexibility tasks, stress-reactivity tasks, and attention control tasks can be analogized to animal paradigms more directly than subjective scales alone. This is one reason reward-bias tasks and reversal learning paradigms are so strategically important: they give a cross-species language for the same functional domain.

The ideal bridging package is therefore multi-modal: receptor occupancy to anchor exposure, EEG/imaging to capture brain-state change, peripheral markers to scale, and tasks to measure function. None of these is definitive alone; together they can support or falsify the hypothesis that plasticity and durable functional change are being produced with minimal acute phenomenology.

8.2 Trial design issues: expectation effects, blinding, and functional outcomes

Psychedelic trials have a unique problem: expectancy is often inseparable from experience. For classic psychedelics, participants know they received the active drug, and that knowledge can change outcomes. That doesn't invalidate psychedelic therapy, but it complicates claims about pharmacology vs context.

For "therapeutic without trip," expectancy becomes both a risk and an opportunity. It is a risk because public narrative can still shape outcomes: people may expect "a psychedelic-like reset" even if the drug is designed not to produce it. It is an opportunity because experience-

minimized compounds may allow better blinding, which improves interpretability of efficacy signals.

Blinding and active controls.
A core advantage of experience-minimized analogues is that they can, in principle, enable double-blind designs with credible placebos or active controls that mimic peripheral sensations without producing strong psychoactive signatures. But this only works if the compound truly produces minimal acute phenomenology at therapeutic exposures. If subtle effects are present and participants detect them, blinding degrades again.

Therefore, early-phase trials should explicitly measure blinding integrity: ask participants and staff to guess assignment and quantify confidence. If blinding holds, efficacy claims become much stronger relative to the classic psychedelic trial landscape.

Expectation management and participant selection.
Trials should either minimize expectancy by neutral framing or measure expectancy explicitly and incorporate it into analysis. In experience-minimized programs, it may also be prudent to include participants who are not recruited from “psychedelic enthusiast” populations, because those populations have strong prior narratives that can contaminate interpretation.

Functional outcomes vs symptom scales.
Symptom scales are important, but the “therapeutic without trip” program should prioritize function—because the program is explicitly claiming practical clinical value, scalability, and deliverability. Functional outcomes include:

- Work and daily living performance
- Social and interpersonal functioning
- Cognitive performance in task-based metrics
- Sustained reductions in relapse behavior or craving-related choices in addiction contexts
- Durable changes in reward learning and decision making

If a compound improves symptom scales but does not improve function or durability, it will be hard to justify as a transformative category.

Durability and dosing schedule.
A core claim in the psychoplastogen narrative is durability beyond acute drug presence. Trials should be designed to test that, not assume it. That means follow-ups at multiple time points and explicit comparison of single-dose vs repeated-dose regimens if the mechanistic hypothesis allows it. It also means measuring whether benefits persist without ongoing intensive psychotherapy, because scalability is part of the design goal.

8.3 Indication matching: depression vs addiction vs cognition

Not all “therapeutic without trip” candidates should be tested in the same indications. Different scaffolds and signaling profiles may naturally align with different clinical needs.

Depression.
Depression is the most common initial target because of unmet need and because several preclinical paradigms and human task measures exist. In depression, the key translation targets include:

- Rapid relief of core mood symptoms
- Reduction of anhedonia and rumination
- Restoration of reward sensitivity and learning
- Improvement in cognitive flexibility and stress resilience

Experience-minimized compounds may be especially attractive for depression if they can deliver meaningful benefit without requiring long supervised sessions and if they can be integrated into standard outpatient care.

Addiction.
Addiction translation is complex because outcomes are behavioral and environmental. The value of a psychoplastogen-like intervention in addiction may lie in reducing cue reactivity, reopening motivational flexibility, and supporting sustained behavior change. But addiction trials must address confounds: social instability, polysubstance use, and the role of community support.

Experience-minimized compounds might be valuable here because they can be administered more routinely and to populations for whom intense psychedelic experiences may be risky or unacceptable. However, addiction also raises an abuse-liability concern: a compound that improves mood or reduces distress could be misused. Translation must therefore include careful monitoring of reinforcement signals and misuse patterns.

Cognition and cognitive symptoms.
Cognitive flexibility, executive function, and reward learning deficits are increasingly recognized as central in depression and addiction, not separate from them. Some analogues may be better positioned as “cognitive repair” agents, especially if they show strong reversal learning or reward task rescue under stress.

For cognition-focused development, the translational strategy should emphasize task batteries and objective performance measures, with imaging and EEG as mechanistic supports. The field

should be cautious about over-claiming “cognitive enhancement”; the more defensible target is “restoration under deficit” and “stress-related cognitive repair.”

Indication matching should therefore be a design decision, not a marketing decision. The same compound can be tested in multiple domains eventually, but early development should choose an indication where the preclinical signal and the mechanistic rationale align, and where objective outcomes can be measured.

8.4 The regulatory reality: defining benefit without relying on subjective “trip” endpoints

Regulators are not primarily interested in philosophical debates about whether the trip is necessary. They care about measurable benefit, acceptable risk, and clear labeling that can be operationalized in clinical settings. For classic psychedelic therapies, subjective experience can be part of the treatment model. For “therapeutic without trip,” the regulatory framing shifts: benefit must be defined in terms of symptom improvement and function, not in terms of induced experience.

This imposes three requirements.

First, clarify the intended use model. Is the compound a single-dose intervention with durable benefit? Is it a short course? Is it maintenance dosing? The safety and efficacy burden differs dramatically across these models, and regulators will evaluate risk in that context.

Second, show that benefit is not merely acute mood elevation. Experience-minimized compounds still risk being interpreted as “rapid mood boosters” unless durability and functional outcomes are demonstrated. This is where the psychoplastogen framing helps: if benefits persist beyond acute exposure and correlate with objective measures, the case strengthens.

Third, address safety in a scalability-aware way. A compound intended for broad outpatient use faces a higher tolerability bar than one intended for supervised administration. Cardiovascular risk, seizure risk, drug–drug interactions, and long-term liabilities (including any 5-HT2B-related concerns under repeated exposure) must be confronted early. The absence of dramatic subjective effects does not reduce regulatory scrutiny; it can increase it, because outpatient scaling increases total exposure across populations.

Finally, regulators will demand precision in claims. A label cannot say “non-hallucinogenic” unless human evidence supports it, and even then it would likely be framed more carefully (for example, “did not produce clinically meaningful psychedelic effects in controlled studies at therapeutic doses”). The field should plan for that language from the beginning rather than retrofitting it after public narrative has hardened.

The deeper lesson is that translation is partly measurement engineering and partly rhetoric discipline. If “therapeutic without trip” is to become a credible category, it must build a

coherent bridge from receptor and proxy data to human claims that regulators can evaluate and clinicians can deliver. The bridge is not a metaphor. It is a set of measurements, designs, and definitions that prevent the science from being lost in its own excitement.

9. Risks, Ethics, and Public Language

The "therapeutic without trip" program is not only a scientific project. It is also a social project that lives inside media cycles, investment incentives, regulatory politics, and the public's unstable relationship with psychedelics. That social context matters because it shapes what gets funded, how quickly claims spread, what patients expect, and what regulators and clinicians feel pressured to accept or reject. In a field where the public narrative is already dramatic—miracles, awakenings, stigma, moral panic—language becomes a mechanism. It can either stabilize trust or destroy it.

This section treats risk and ethics in a practical sense: how to avoid predictable failure modes that come from overstated claims, misconstrued proxies, and commercialization dynamics. The goal is not to moralize the enterprise. It is to keep it real enough to survive its own momentum.

9.1 Hype control: avoiding the false binary of "hallucinogenic vs therapeutic"

The first ethical risk in this space is conceptual. The public wants a binary: either psychedelics heal because of the trip, or psychedelics heal despite the trip. That binary is satisfying, but it is almost certainly false.

Therapeutic outcomes likely emerge from a composite of factors: receptor pharmacology, circuit dynamics, plasticity windows, learning context, and subjective meaning. For some patients and indications, subjective experience may be essential. For others, it may be optional. For still others, it may be actively harmful. The correct landscape is pluralistic, not binary.

Hype tends to collapse pluralism into slogans. "Non-hallucinogenic psychedelic" is one such slogan: it promises a magical divorce between the dramatic and the useful. The field should resist this framing for two reasons.

First, it encourages premature certainty. In early stages, the evidence usually supports "reduced hallucinogen-like proxies in animals," not "no psychedelic effects in humans." When a slogan outruns evidence, it sets up a future correction that will look like scandal even if the science was simply early.

Second, it implies that hallucinogenic and therapeutic are opposites, as if the hallucinogenic aspect is merely a liability. That can be ethically corrosive. It can delegitimize patients whose

benefit is tied to experience-centered therapy and can split the field into ideological camps rather than methodological camps.

Hype control, in this context, means teaching the public a more accurate shape: not a binary but a trade space. A compound can be more or less perceptually disruptive, more or less durable, more or less scalable, more or less safe for vulnerable populations. The ethical communication challenge is to replace slogans with gradients while still telling a compelling story.

A useful discipline is to treat "therapeutic without trip" as a hypothesis with measurable predictions rather than as a marketing category. The prediction is not "no trip." The prediction is "durable benefit with reduced acute phenomenology and improved scalability." That is testable, and it can be stated without claiming more than is known.

9.2 Misuse risk and access: what changes when compounds feel "ordinary"

There is a paradox at the heart of experience-minimized design: making a compound feel more ordinary may increase access and reduce stigma, but it may also change misuse patterns.

Classic psychedelics have an unusual abuse profile. Many are not typically used compulsively in the way stimulants or opioids are, and their intensity can act as a natural brake. If a compound is engineered to be experience-minimized, it might lose that brake. Even if it is not euphoric, it might become a tool for emotional regulation, productivity, or "quiet optimization" in ways that create new kinds of dependence.

Misuse risk here is not only about addiction. It includes:

- Casual use without medical supervision because the drug "doesn't feel like a psychedelic."
- Overuse driven by the belief that "if it's non-hallucinogenic it must be safe."
- Use in vulnerable populations (bipolar vulnerability, psychosis risk) under the false assumption of zero perceptual liability.
- Use as a lifestyle enhancer in high-pressure environments, potentially creating coercive norms ("everyone at work is taking it").
- Off-label stacking with other psychiatric medications or substances, creating unpredictable interactions.

Access also changes when the compound feels ordinary. If a drug can be administered like standard pharmacotherapy, more clinicians may be willing to prescribe it. That is a public good, but it amplifies the importance of accurate labeling and clinician education. The ethical stance

cannot be “we removed the trip, therefore the risk is gone.” The ethical stance must be “we changed the risk shape.”

That risk shape depends on dosing model. A one-time supervised intervention creates different misuse dynamics than a take-home outpatient medication. If a compound becomes a routine prescription, the field must plan for the realities of real-world prescribing: incomplete monitoring, inconsistent follow-up, uneven clinician training, and the tendency for medications to drift into broad use beyond their original intended population.

Therefore, experience-minimized design should be paired with “experience-minimized overconfidence.” Making a drug feel ordinary increases the need for disciplined guardrails, not the opposite.

9.3 Equity and commercialization: patents, startups, and the ownership of scaffolds

The third ethical fault line is ownership. Psychedelic-inspired therapeutics sit at a contentious intersection of public heritage and private property. Many scaffolds originate in nature or in long-standing scientific literature. Yet development requires capital-intensive work: synthesis, profiling, toxicology, manufacturing, clinical trials. Commercialization is not inherently unethical; it is often the only path to approved medicines. But commercialization changes incentives, and incentives shape language and access.

Patents and exclusivity are a double-edged instrument. On one hand, they can finance development. On the other hand, they can narrow access and distort research priorities toward “patentable novelty” rather than “clinical necessity.” Some design moves may be chosen because they are defensible IP rather than because they are the best therapeutic solution. In the “therapeutic without trip” space, where minor edits can produce major functional shifts, the temptation toward IP-driven micro-novelty is obvious.

Equity concerns arise in at least three forms:

Pricing and access.
If a compound is designed to be scalable but priced as a luxury, the social promise collapses. This is especially sharp in depression and addiction, where burden is often concentrated in lower-resource communities.

Trial representation.
If early human trials are run primarily in narrow, affluent cohorts, the translation will be biased and later real-world use will reveal surprises. Experience-minimized compounds may allow broader recruitment than classic psychedelics; the field should use that advantage intentionally.

Ownership of scaffolds and cultural narratives.
There is a deeper unease when private entities claim ownership over small edits to molecules

that feel culturally “shared,” especially when those molecules have indigenous histories or long-standing public scientific histories. Even when legally permissible, this can produce public backlash that harms trust.

A practical ethical posture here is transparency and reciprocity. Be explicit about what is being patented (a specific analogue, a specific synthesis route, a specific use case), be explicit about what remains in the public domain, and commit—where possible—to access models that match the public narrative of medical benefit. If a company markets “we’re making psychedelic healing accessible,” but structures pricing to maximize exclusivity, the gap between story and reality becomes an ethical wound.

9.4 A disciplined vocabulary for journalists, investors, and clinicians

Because language is the primary vector for hype and misunderstanding, the field needs a shared vocabulary that different audiences can use without inadvertently lying. Below is a practical lexicon—terms and recommended usage—that can reduce misinterpretation.

Recommended terms and what they mean:

“Psychedelic-inspired therapeutic”
Use when a compound was motivated by psychedelic biology or scaffolds but may not be a classic psychedelic in human experience.

“Analogue”
Use to describe structural lineage, not subjective similarity. Always specify what scaffold family it is analogous to.

“Reduced hallucinogen-like proxy activity”
Use for preclinical claims. Specify which proxy (for example, HTR) and the exposure context.

“Experience-minimized”
Reserve for human evidence. Even then, define the measurement: subjective scales, clinician-rated effects, functional impairment, etc.

“Psychoplastogen / neuroplastogen”
Use as functional categories. Pair with explicit evidence: plasticity markers and durable behavioral effects, not just receptor binding.

“Biased agonism / functional selectivity”
Use only with context: assay system, pathways measured, dose and time dependence. Avoid using it as a synonym for “safer.”

Phrases to avoid or qualify heavily:

“Non-hallucinogenic psychedelic”
If used at all, qualify as “non-hallucinogenic in preclinical proxy assays” and explain that this does not prove absence of altered experience in humans.

“Trip-free”
Avoid as a blanket claim. If used, it must be bound to dose, population, and measurement, and even then it is usually too categorical.

“Same benefits without the trip”
Avoid unless clinical evidence supports it. More responsible: “aims to preserve plasticity-linked and therapeutic-like effects while reducing acute psychedelic effects.”

A simple audience-specific discipline:

For journalists:
Describe what was measured. Name the proxy. Avoid translating proxies into human experience statements. Use “candidate” and “early evidence” language.

For investors:
Treat proxy reduction as de-risking, not de-risked. Ask for the evidence ladder: receptor → plasticity → behavior under deficit → safety → human phenomenology → clinical efficacy.

For clinicians:
Demand human data before accepting “experience-minimized” claims. Focus on delivery model, contraindications, drug–drug interactions, and real-world monitoring burdens.

A final ethical point: disciplined vocabulary is not a marketing disadvantage. It is a long-term trust advantage. Psychedelic-inspired medicine is one of the few domains where overstatement can trigger not just disappointment but cultural backlash and regulatory chill. The field’s best defense is precision: say what you did, say what it means, and say what it does not yet mean.

11. References

11.1 Primary papers on JRT and LSD-analogue design

1. Tuck, J. R., et al. *Molecular design of a therapeutic LSD analogue with reduced hallucinogenic potential.* Proceedings of the National Academy of Sciences (PNAS). 2025. doi:10.1073/pnas.2416106122. PubMed+1

2. Karst, M., Halpern, J. H., Bernateck, M., Passie, T. *The non-hallucinogen 2-bromo-lysergic acid diethylamide as preventative treatment for cluster headache: an open, non-randomized case series.* Cephalalgia. 2010;30(9):1140–1144. doi:10.1177/0333102410363490. PubMed

3. Wacker, D., et al. *Crystal Structure of an LSD-Bound Human Serotonin Receptor.* Cell. 2017;168(3):377–389.e12. doi:10.1016/j.cell.2016.12.033. PubMed

4. Nichols, D. E., Frescas, S., Marona-Lewicka, D., Kurrasch-Orbaugh, D. M. *Lysergamides of isomeric 2,4-dimethylazetidines map the binding orientation of the diethylamide moiety in the potent hallucinogenic agent N,N-diethyllysergamide (LSD).* Journal of Medicinal Chemistry. 2002. doi:10.1021/jm020153s. PubMed

5. Brandt, S. D., Kavanagh, P. V., Nichols, D. E. *Lysergamides and their analogues: a review.* (Select review articles and updates across the "return of the lysergamides" literature). PMC

11.2 Key neuroplastogen and psychoplastogen discovery papers

1. Olson, D. E. *Psychoplastogens: A Promising Class of Plasticity-Promoting Neurotherapeutics.* Journal of Experimental Neuroscience. 2018. doi:10.1177/1179069518800508. PubMed

2. Ly, C., et al. *Psychedelics Promote Structural and Functional Neural Plasticity.* Cell Reports. 2018. (Key "psychoplastogen"/rapid structural plasticity foundational paper). PubMed

3. Cameron, L. P., et al. *A non-hallucinogenic psychedelic analogue with therapeutic potential.* Nature. 2021;589(7842):474–479. doi:10.1038/s41586-020-3008-z. PubMed

4. Dong, C., Ly, C., Dunlap, L. E., et al. *Psychedelic-Inspired Drug Discovery Using an Engineered Biosensor.* Cell. 2021;184(10):2779–2792.e18. doi:10.1016/j.cell.2021.03.043. PMC

5. Lu, J., et al. *An analog of psychedelics restores functional neural circuits disrupted by unpredictable stress.* Molecular Psychiatry. 2021. doi:10.1038/s41380-021-01061-9. Nature

6. Ly, C., et al. *Transient stimulation with psychoplastogens is sufficient to initiate neuronal growth.* ACS Pharmacology & Translational Science. 2020. doi:10.1021/acsptsci.0c00065. ScienceDirect

7. Salfati, E. L., et al. *Zalsupindole is a non-hallucinogenic 5-HT2A agonist with antidepressant-like activity.* ACS Chemical Neuroscience. 2025. doi:10.1021/acschemneuro.5c00667. pubs.acs.org

8. Bahlouli, W., et al. *Non-hallucinogenic psychedelics for mood and anxiety disorders: A systematic review.* Journal of Affective Disorders. 2025. (Useful map of the preclinical evidence base and endpoints used).

11.3 Receptor signaling, biased agonism, and 5-HT2A mechanistic literature

1. Nichols, D. E. *Psychedelics.* Pharmacological Reviews. 2016;68(2):264–355. doi:10.1124/pr.115.011478. PubMed

2. Vollenweider, F. X., Preller, K. H. *Psychedelic drugs: neurobiology and potential for treatment of psychiatric disorders.* Nature Reviews Neuroscience. 2020.

3. Urban, J. D., et al. *Functional selectivity and classical concepts of quantitative pharmacology.* Journal of Pharmacology and Experimental Therapeutics. 2007.

4. Kenakin, T. *Biased Receptor Signaling in Drug Discovery.* Pharmacological Reviews. 2019;71(2):267–315. doi:10.1124/pr.118.016790.

5. González-Maeso, J., et al. *Hallucinogens recruit specific cortical 5-HT2A receptor-mediated signaling pathways to affect behavior.* Neuron. 2007.

6. Cao, C., et al. *Structure-based discovery of nonhallucinogenic psychedelic analogs.* Science. 2022.

7. Kim, K., et al. *Structure of a hallucinogen-activated Gq-coupled 5-HT2A serotonin receptor.* Cell. 2020.

8. Wallach, J., et al. *Identification of hallucinogenic and non-hallucinogenic psychedelic analogs and their receptor signaling profiles.* Nature Communications. 2023.

9. Gumpper, R. H., et al. *Structural diversity of psychedelic drug actions revealed by seven high-resolution cryo-EM structures of 5-HT2A receptor signaling complexes.* Nature Communications. 2025.

10. Wacker, D., et al. *Crystal Structure of an LSD-Bound Human Serotonin Receptor.* Cell. 2017;168(3):377–389.e12. doi:10.1016/j.cell.2016.12.033.

11.4 Preclinical proxy assay validation and limitations

1. Halberstadt, A. L., Geyer, M. A. *The head-twitch response as a behavioral assay for 5-HT2A receptor activation and hallucinogen-like activity.* (Foundational validation/interpretation work across multiple papers and reviews).
2. Canal, C. E., Morgan, D. *Head-twitch response quantification methods and pharmacological validation.* (Early automation/quantification and methodological grounding).
3. de la Fuente Revenga, M., et al. *Fully automated head-twitch detection (video/ML or magnetometer-based) and validation against known 5-HT2A agonists/antagonists.* Scientific Reports. 2019.
4. de la Fuente Revenga, M., et al. *Improved automated HTR detection and cross-lab reproducibility considerations.* 2020.
5. Dong, C., et al. *Psychedelic-Inspired Drug Discovery Using an Engineered Biosensor.* (Illustrates coupling of in vitro "hallucinogen-likeness" proxies + HTR + antidepressant-like behavioral assays; also documents failure modes where behavior shifts without HTR). Cell. 2021. doi:10.1016/j.cell.2021.03.043.
6. Bahlouli, W., et al. *Non-hallucinogenic psychedelics for mood and anxiety disorders: A systematic review.* (Summarizes which animal endpoints are being used as "non-hallucinogenic" proxies and flags evidence limitations). Journal of Affective Disorders. 2025.

11.5 Clinical and regulatory perspectives on psychedelic-inspired therapeutics

1. U.S. Food and Drug Administration (FDA). *Psychedelic Drugs: Considerations for Clinical Investigations.* Draft Guidance for Industry. June 2023.
2. Johnson, M. W., Richards, W. A., Griffiths, R. R. *Human hallucinogen research: guidelines for safety.* Journal of Psychopharmacology. 2008.
3. Carhart-Harris, R. L., et al. *Trial of psilocybin versus escitalopram for depression.* New England Journal of Medicine. 2021.
4. Goodwin, G. M., et al. *Single-dose psilocybin for treatment-resistant depression.* New England Journal of Medicine. 2022.

5. Davis, A. K., et al. *Effects of psilocybin-assisted therapy on major depressive disorder.* JAMA Psychiatry. 2021.

6. Mitchell, J. M., et al. *MDMA-assisted therapy for severe PTSD: a randomized, double-blind, placebo-controlled phase 3 study.* Nature Medicine. 2021. doi:10.1038/s41591-021-01336-3.

7. Vollenweider, F. X., Preller, K. H. *Psychedelic drugs: neurobiology and potential for treatment of psychiatric disorders.* Nature Reviews Neuroscience. 2020. (High-level clinical translation framing, endpoints, and open questions).

The author has published over 4,800 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.